Trusted Multi-Net

Microsoft Solutions for Multi-Security Domain Access

Sean Finnegan
Bill Neumann

Microsoft Windows

UNCLASSIFIED

Trusted Multi-Net: Typhon XP

Using Virtual Machines to Access Multiple Security Domains

Sean Finnegan
Security Program Manager
Microsoft Federal
seanfi@microsoft.com

Microsoft Windows

Background

- Typhon is a Microsoft Solution Offering
  - Not a new product
  - No new code
  - Entirely COTS
- Provides user access to multiple security domains from one workstation
  - Reduce space, weight, and power requirements
  - Support existing applications
  - Support simultaneous sessions with quick switching
  - Leverage local computing power
- Builds on NSA research using VMMs to provide separation
  - Host OS a “black hole” on the network
  - Provide multiple layers of isolation

Demonstration

Typhon XP User Experience

Typhon XP Architecture

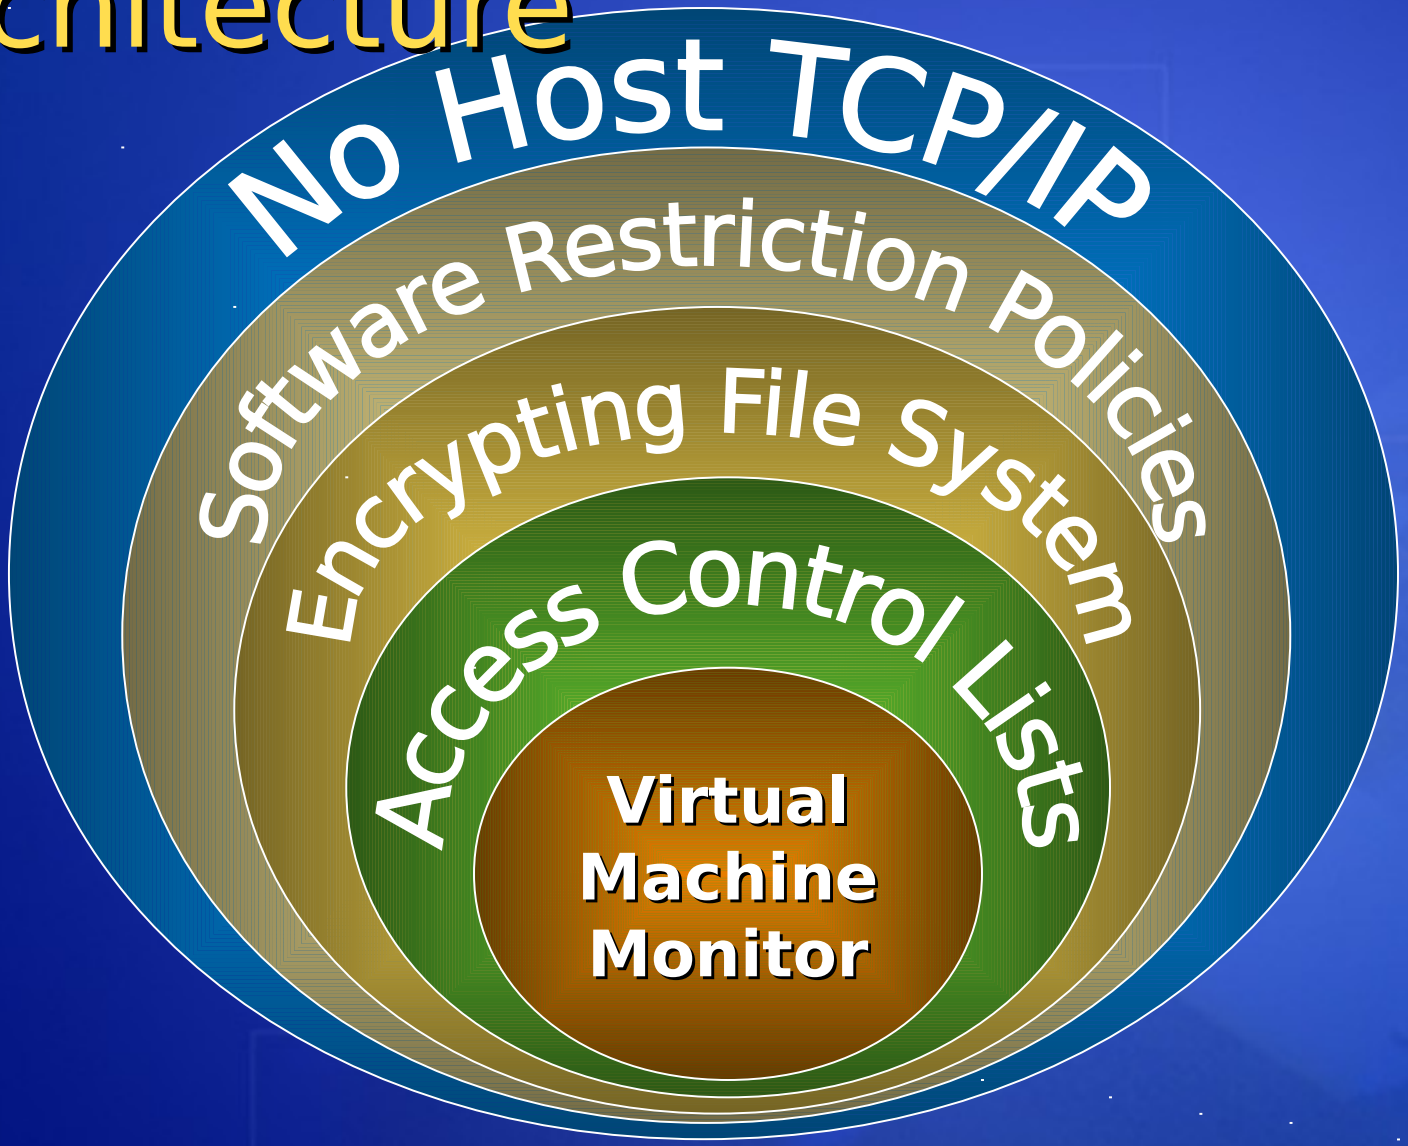

Windows Architecture Basics

- Processes and Tokens
- Access Control Lists
- Terminal Svcs & Fast User Switching
- Windows Network Stack
- Software Restriction Policies
- Encrypting File System

Processes & Tokens

- Security Reference Monitor (SRM) in kernel controls access to objects
- Each process or thread has an associated *access token*
- SRM uses the *access token* to identify the processes security context
- Access token contains:
  - SID (security identifier) of user
  - SIDs of groups the user is a member of
  - Privileges held
  - Other stuff...

Processes & Tokens

- Each process runs in own virtual memory space
  - Unless granted a special privilege one user cannot open another user's process
- Threads can change their security context
  - Impersonation using an authenticated connection (e.g. authenticated RPC or named pipe)
- Processes can be created as other than the current user's identity
  - Using LogonUser or CreateProcessAsUser API
  - Must specify user credentials (userid, password, domain)
  - *Only Winlogon process can do this for accounts with blank passwords*
    - Prevents attacks on user accounts with no password

Access Control Lists

- Named objects have a security descriptor (SD)
  - (e.g. files, dirs, reg keys, processes, named pipes, etc.)
- SD contains owner ID, DACL, and SACL
  - SACL describes conditions to audit
- DACL is a list of SIDs and permissions to the object
  - Can be allowed or denied permissions
- Process or threads request a set of permissions when opening an object
  - SRM walks DACL with user and group SIDS from token until desired access it accumulated, access

Terminal Services & Fast User Switching

- Each Win32 process can have one or more windows
- Windows are bound to a given desktop
  - Windows on the same desktop can send messages to one another
    - e.g. mouse click, keystrokes, etc.
  - Each desktop has own memory heap so no window messages can be sent between them
- Desktops are bound to a window station
  - Logged in user typically has multiple desktops
    - Apps, secure screen saver, secure attn sequence
  - Only one window station is visible on the console and can receive keyboard and mouse input at a time
  - Clipboard is per window station and blocked on

Terminal Services & Fast User Switching

- Window stations are part of a session
  - Session abstraction created for terminal services
  - Anything within same session gets the same kernel GUI state
- Terminal Server adds multiple sessions
  - Session keyboard/video/mouse IO is bound to the glass or a network driver
- Fast User Switching also uses multiple sessions
  - Session keyboard/video/mouse IO is bound to the glass or a null device

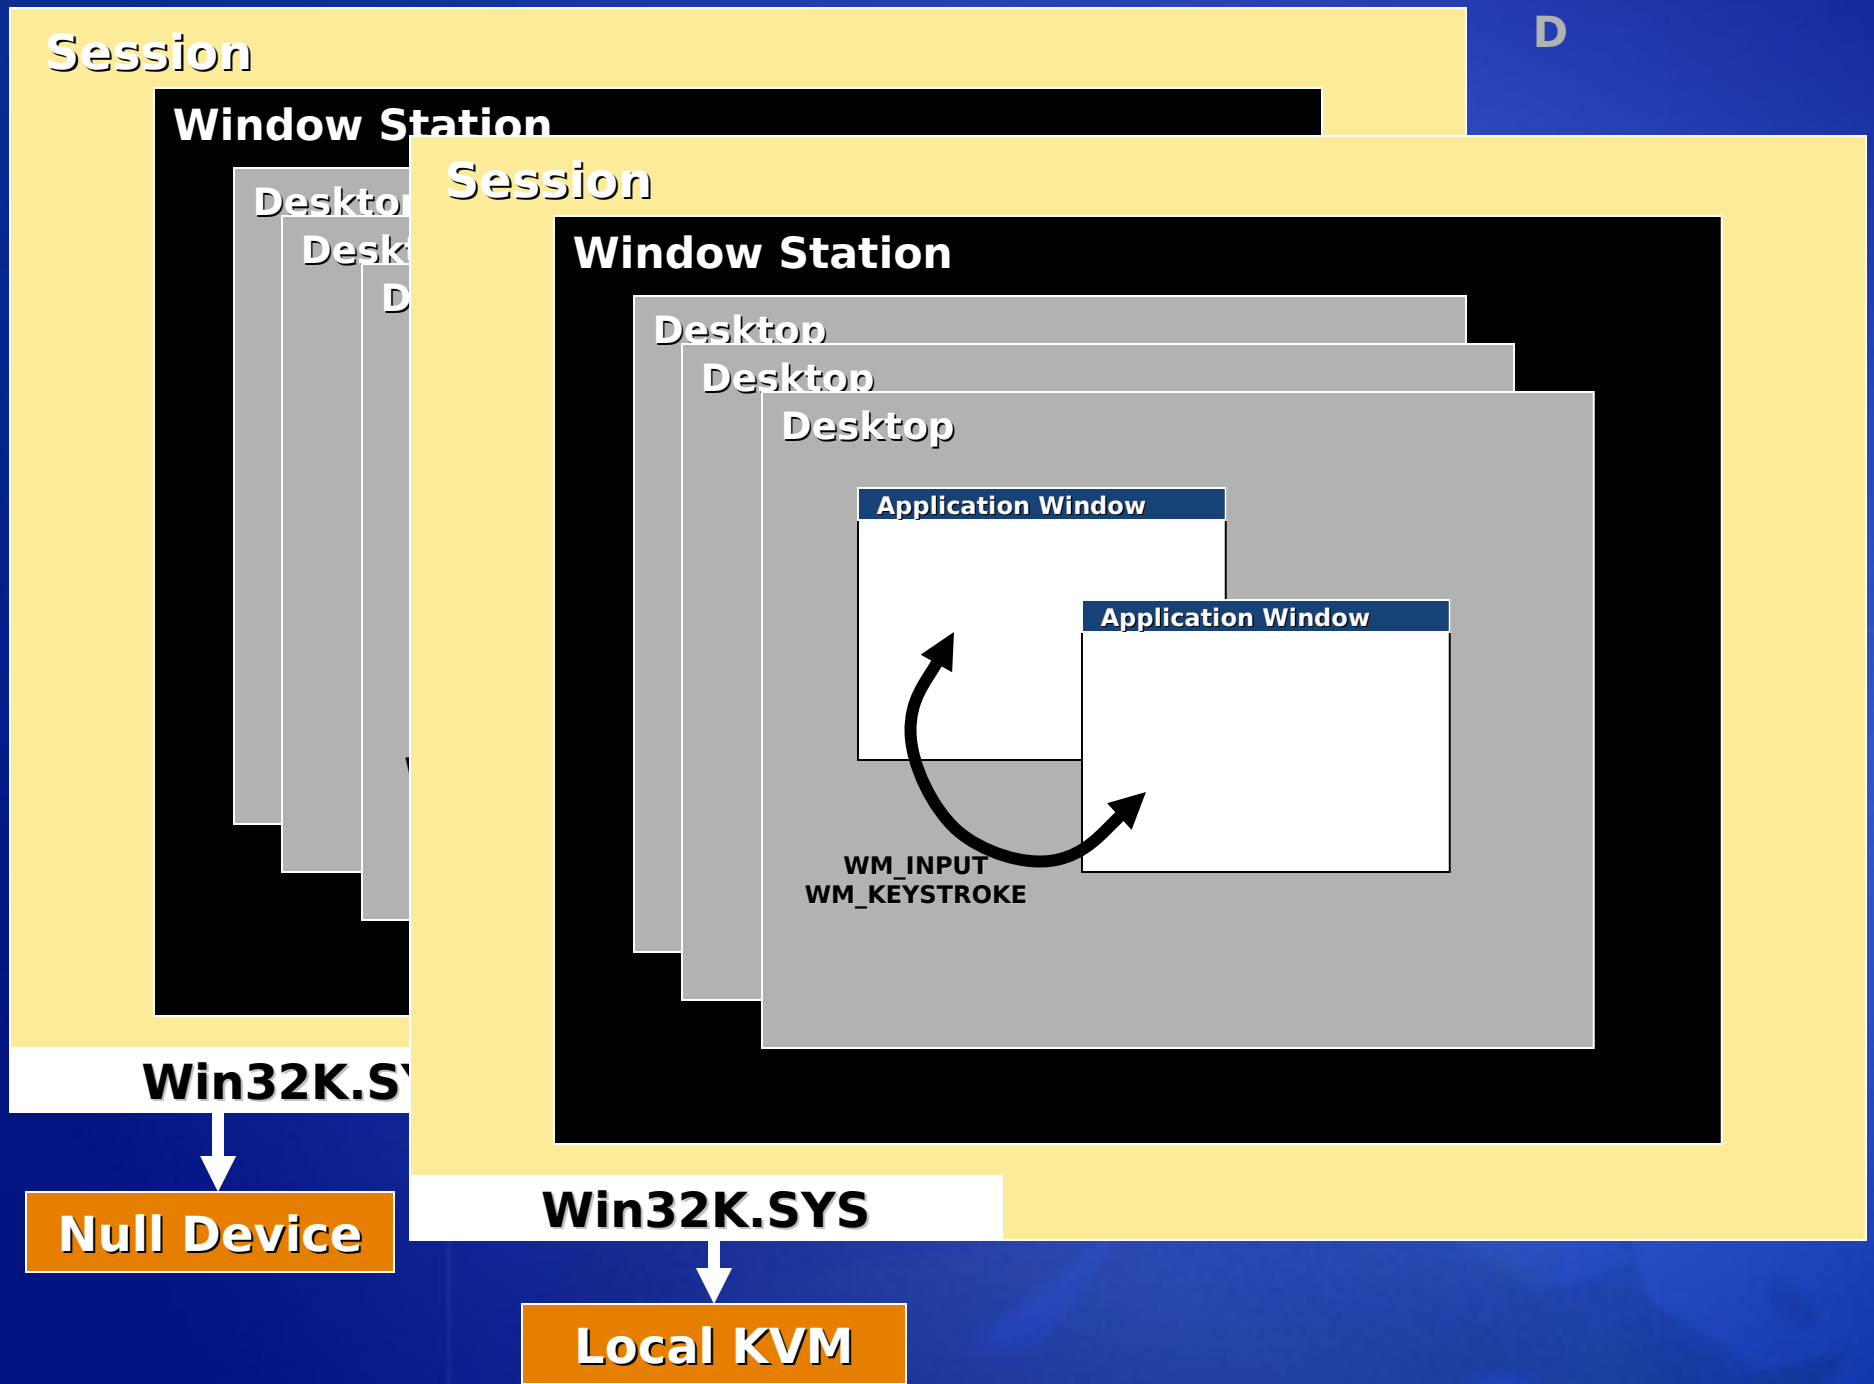

Windows Network Stack

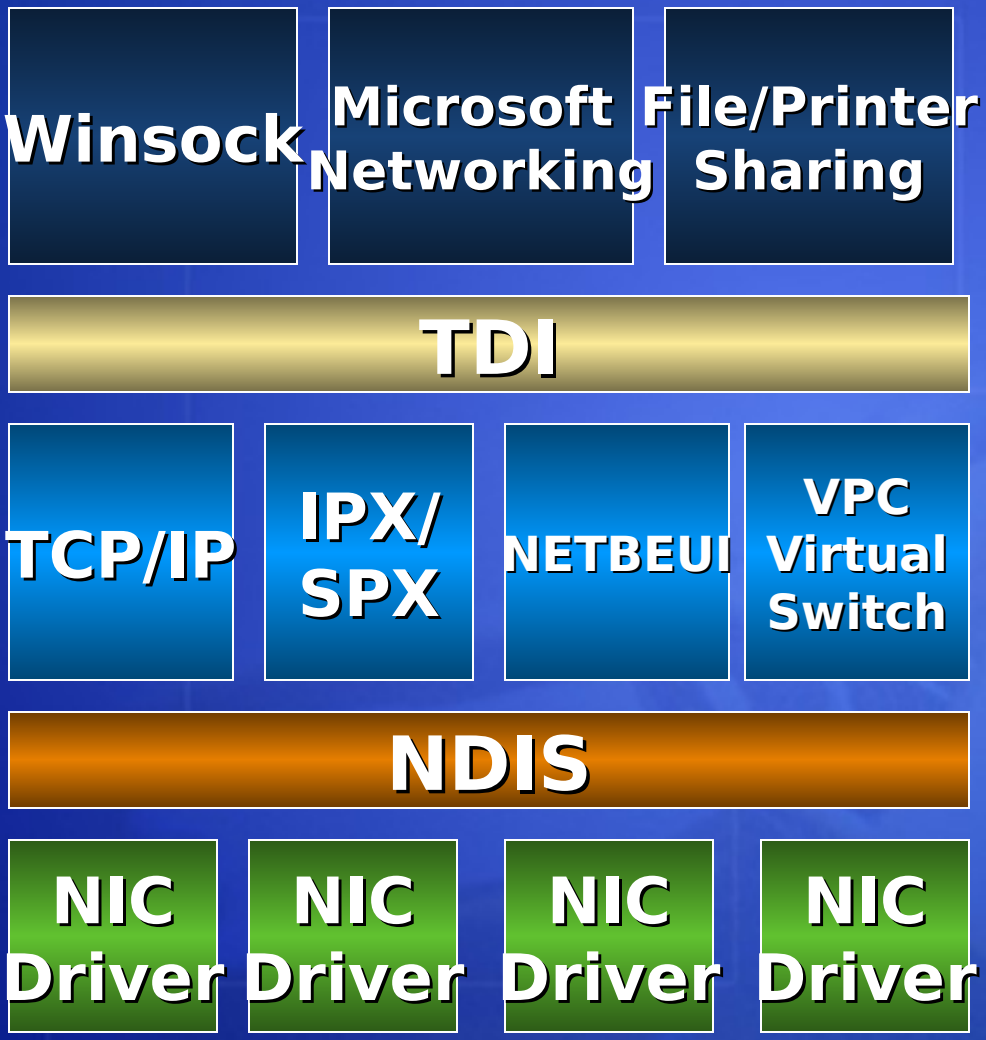

Software Restriction Policies

- Policy based mechanism in Windows XP to restrict program execution
  - Two policies out of the box
    - Trusted: run with full user privileges
    - Untrusted: do not run
  - Two levels out of the box
    - Unrestricted: Run all code except those specified in policy
    - Disallowed: Only run code expressly permitted by policy
    - Can exempt administrators and checks on DLLs
- Code identified by:
  - File/folder paths (e.g. c:\windows\system32)
  - URLs – http paths, UNC paths, etc.
  - Image hashes (e.g. MD5 hash for evil.exe)
  - Publisher certificates (e.g. DoD PKI cert)
- Policy enforced at Win32 CreateProcess and LoadLibrary API for native code

Encrypting File System

- Implemented as a file system filter on Windows starting with Windows 2000
  - Transparent to user and applications
  - Very high performance
- Discretionary to the file and dir level
  - Cannot be applied to the OS files directly
- Windows XP SP1 uses AES-256
  - Private key stored in user profile encrypted by user logon credentials
  - Ability to define recovery agent without private key on the system
- Use on Typhon primarily for additional separation (will discuss later)

File Encryption

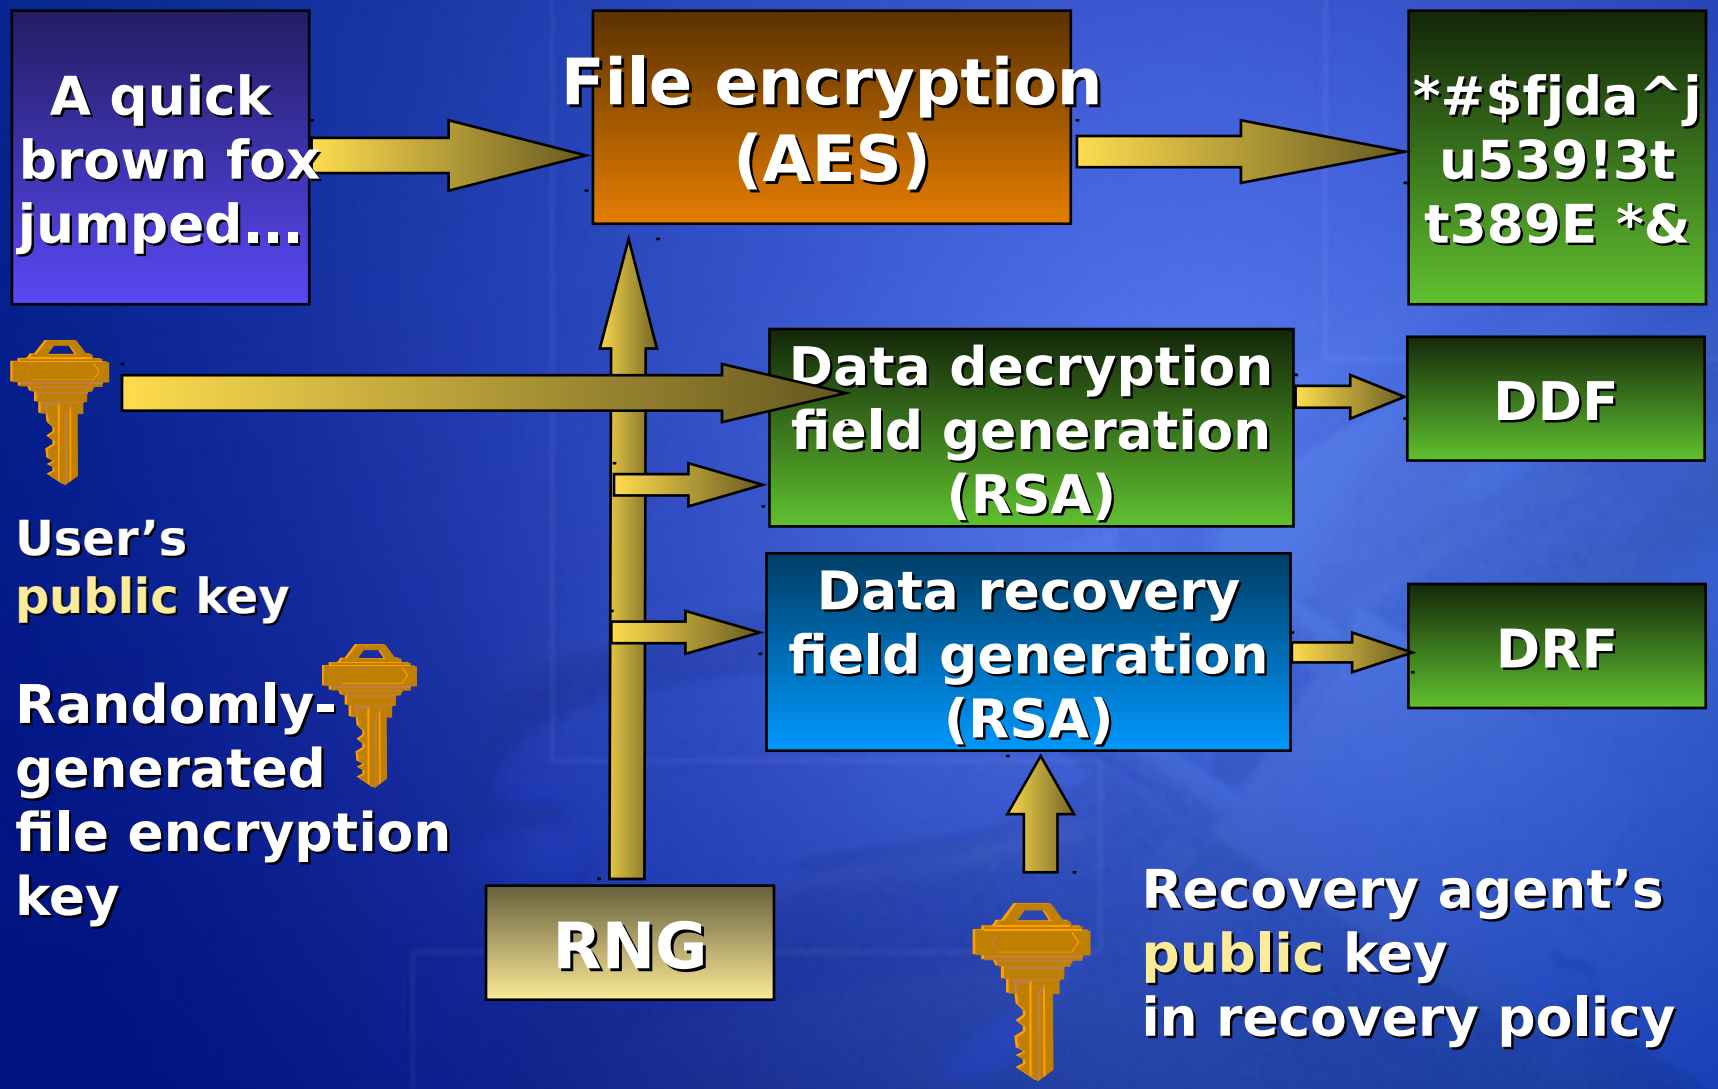

File Decryption

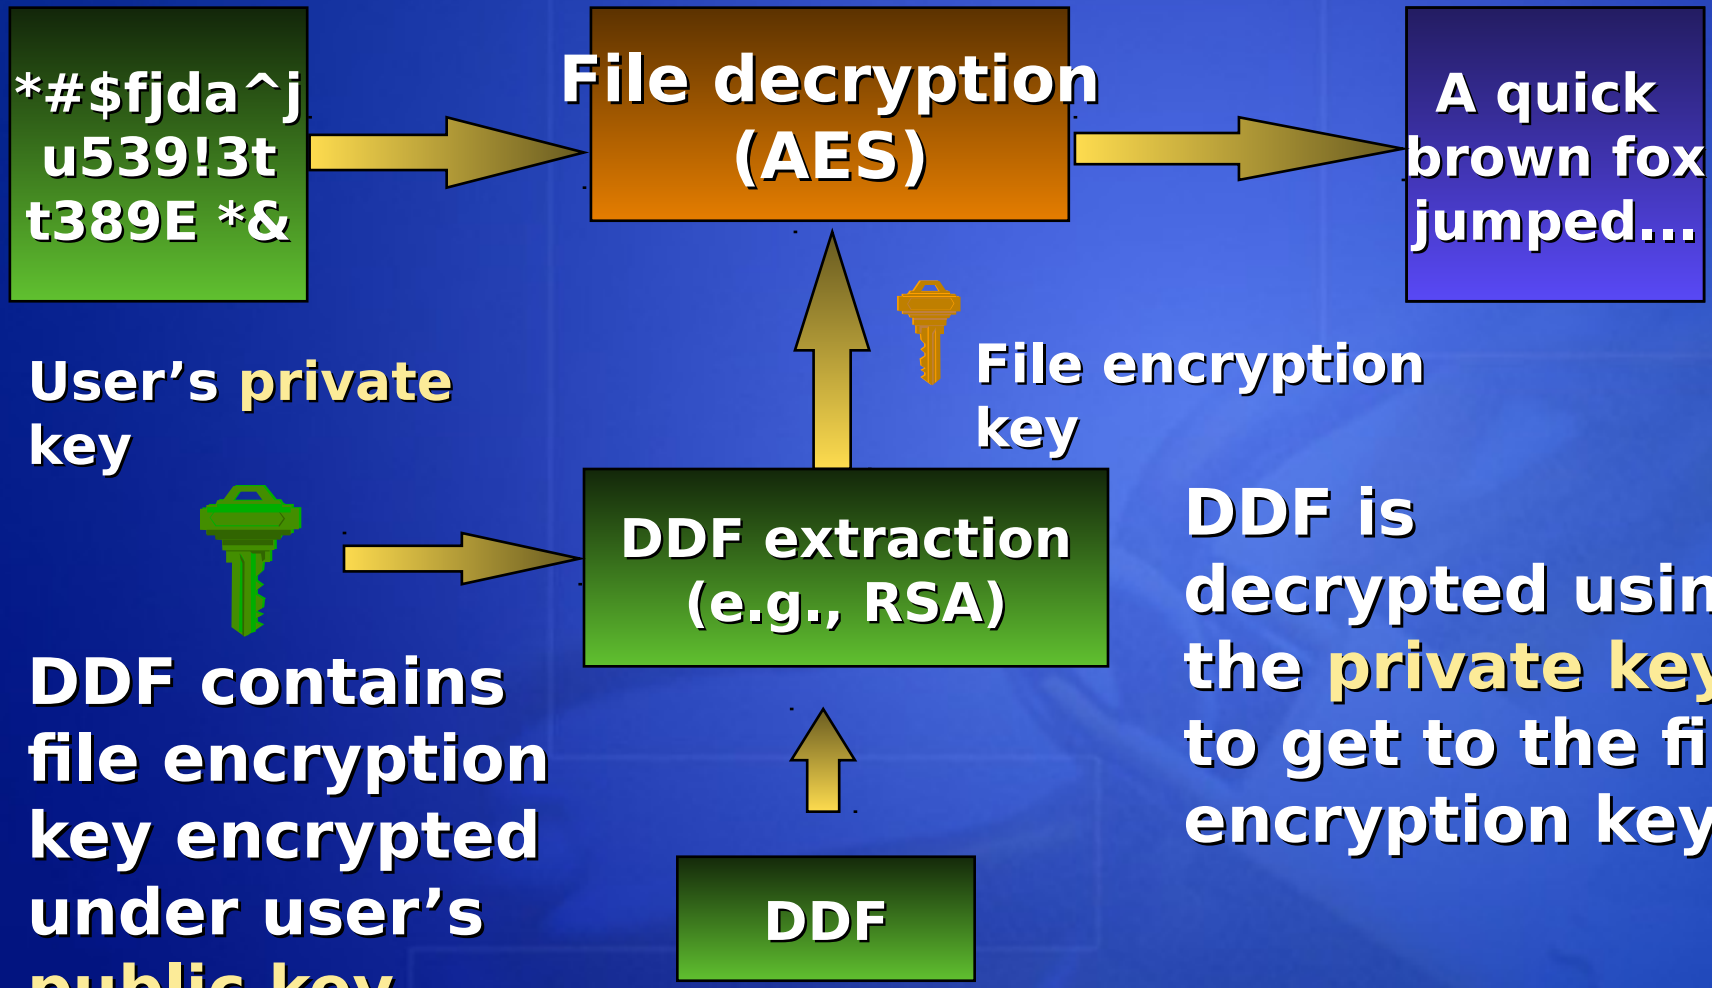

Typhon XP Architecture

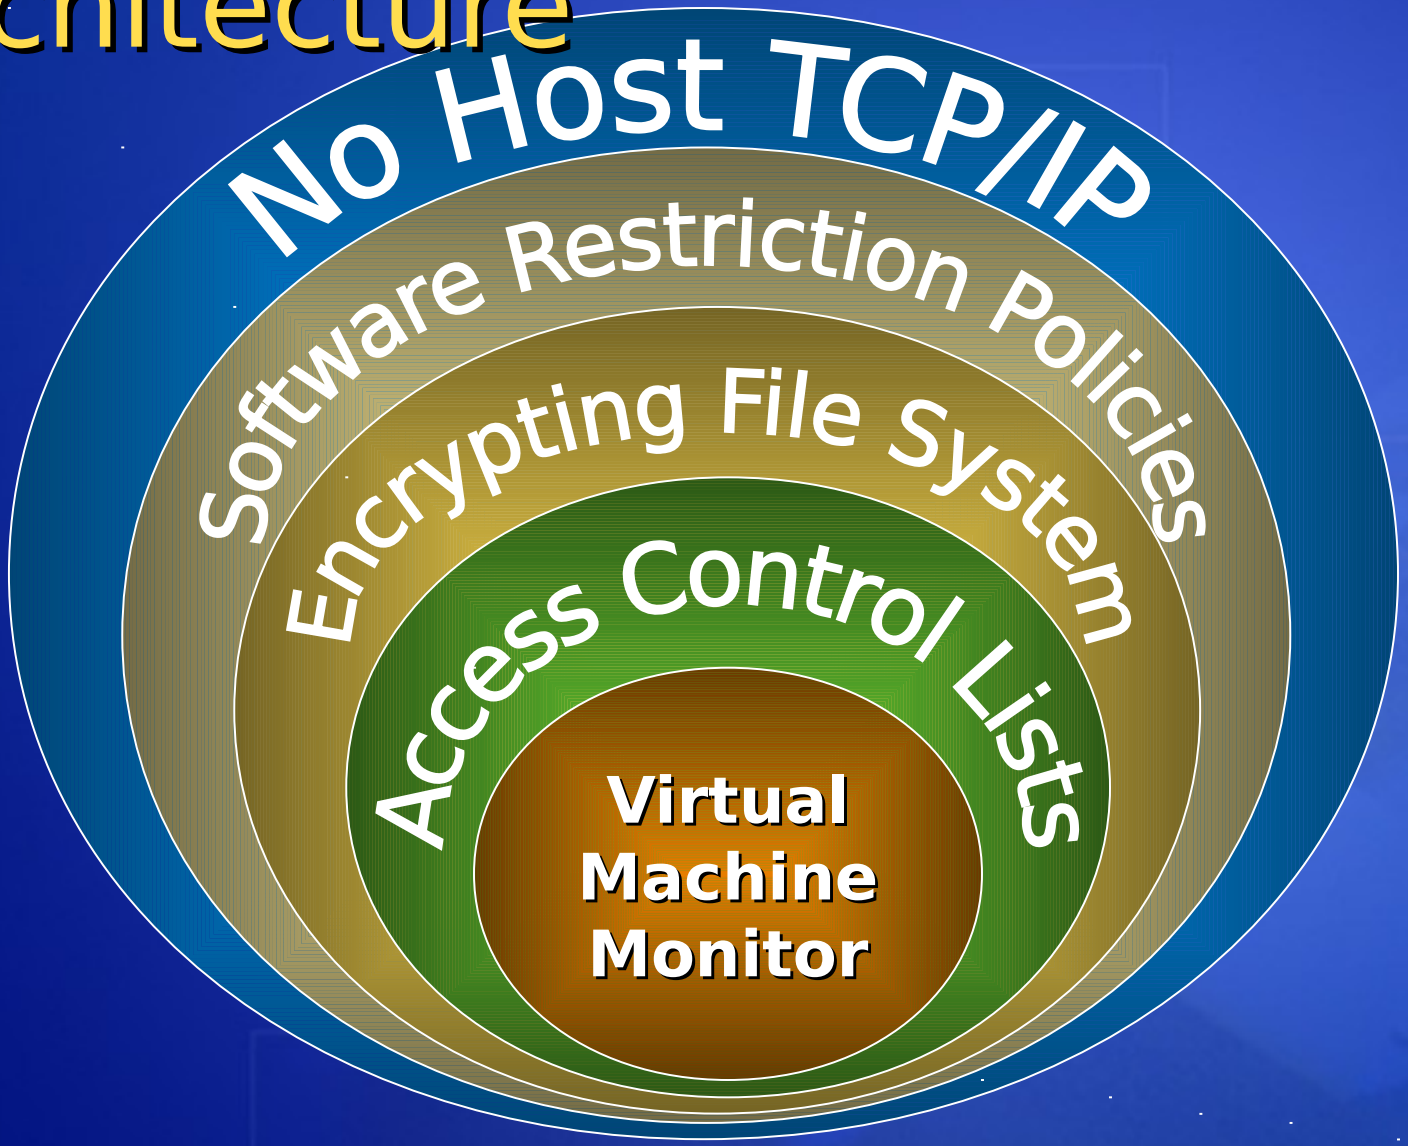

Virtual Machines

- Virtual Machine Monitor software isolates hardware on the host OS from the guest OS
  - Each VM runs as a single process on the host
  - Can use VMWare Workstation or Microsoft Virtual PC 2004
    - VMWare already subject of NSA analysis
- Guest OSs runs within virtual machines
  - Standard commercial OS and apps
  - Each Guest OS can be domain joined and networked including AD for management
- Each VM bound to a single NIC
  - VPN option in development to use single cable to the desktop

Security Contexts

- Each network/level is represented as a local user account on the host
  - E.g. *SIPR*, *Coalition*, *JWICS* user accounts
  - These accounts have no passwords and don't need them
    - Host OS prevents creating process from one account as other
    - VM guest OS still requires user authentication to the network
- Each VM runs as a different local user
  - Separate processes with access tokens
  - Existing process isolation prevent access between these processes

Security Contexts

- Each VM in a different user session
  - Different desktop and window station
  - Different address space
  - One VM window cannot send window messages to others
- Windows XP Fast User Switching used to toggle between levels
  - + L key switches desktops
  - Biometric option to switch between levels
    - Train each finger a different level/user
    - Convenience feature not a security feature
    - Ideal for customers where one user per machine

Selecting Security Domain

Separation

- Each VM uses separate virtual disk file on host OS in own directory
  - Looks to each VM as own hard drive but really just big files on host
  - E.g. D:\low\low.vmdk is low VM “C drive”
- Each disk file ACLed so that only one level user account can access
  - Prevents rogue app at one level from attacking storage of another
- Each VM disk file also encrypted so only the “level” user account can open
  - Secondary isolation with little overhead
  - Provides protection from offline attack on VM disk file if OS disk is not attacked

Host Operating System

- Windows XP Embedded SP1
  - Same binaries as XP Pro but with only what you need to run Typhon
  - Adds some additional capabilities
- 38+ OS services disabled compared with XP Pro
  - While many services have been disabled their files have not removed them from the disk image yet
- User shell set to VMM Application
  - No explorer interface to mess with
- Local account wallpaper set to show current level
- Software Restriction Policies permit only VMM
- Most networking services removed from host OS
  - TCP/IP files still present on current builds but unbound from any NICS
  - VMM bridge or equiv protocol only service bound above NDIS layer

Typhon XP Architecture

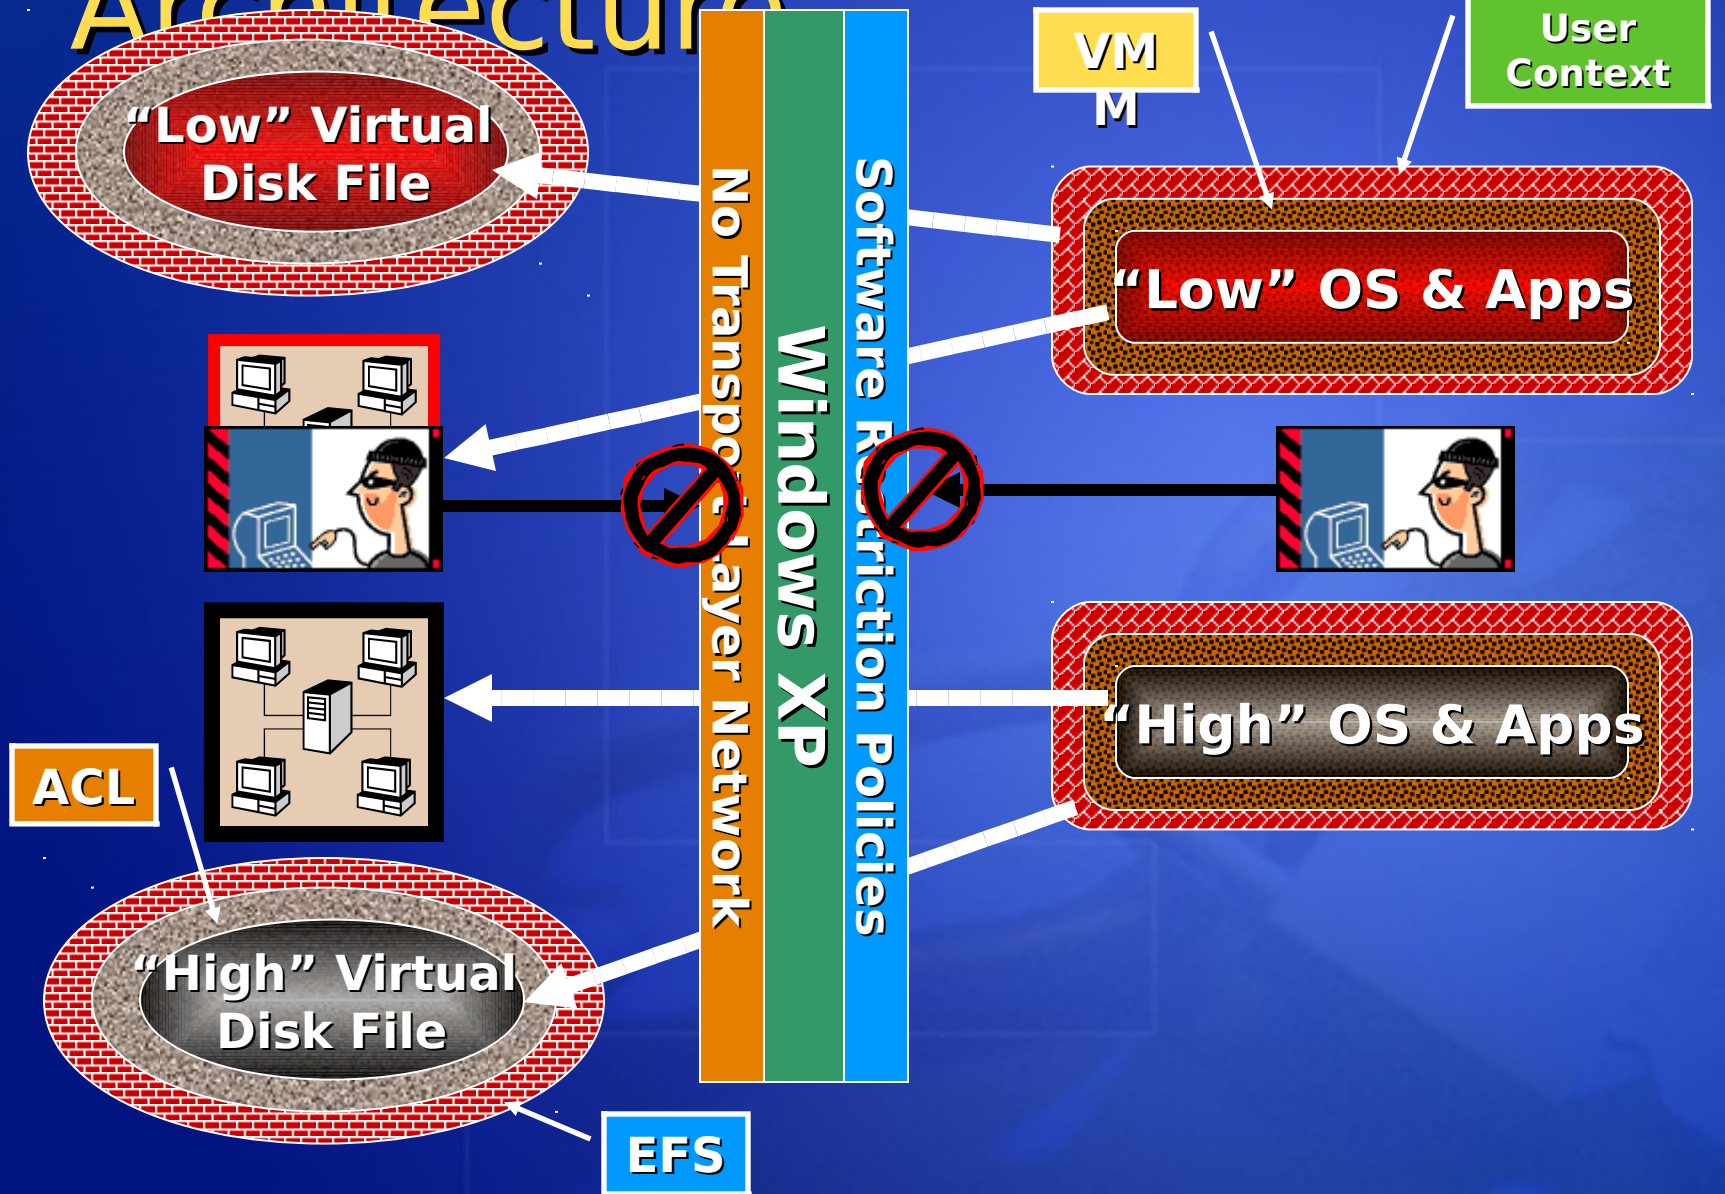

Host Boot Options

- Host OS and VMM are stored in own partition on hard drive or removable media
  - Only changeable state is on data partition for VM disk files
  - Host files can be updated with simple copy
  - Enhanced Write Filter can be used to prevent any writes to host partition
    - EWF is new feature to XP Embedded
    - Redirects all writes to disk to memory
    - Writes appear to succeed but are lost on reboot
- Host OS/VMM can be booted from CD
  - Booting directly from CD is slow until OS files are cached in RAM
  - Require no more memory than hard drive boot

Host Boot Options

- XPe supports copying entire CD to RAM drive and booting from RAM
  - Very fast after initial copy
  - Requires more RAM for OS/VMM RAM drive
  - If host needs to be updated hand out new CDs at the door
- Can also copy OS/VMM image to RAM Drive from server at boot time
  - Uses PXE or BXP boot
  - Update host OS image centrally
  - Once machine is booted it can function through network outages
  - Only data left on machine at power-off is AES-256 encrypted virtual disk files

Data Sharing

- Data sharing between levels at the host was not an original design goal
  - Easier to control data flow through fixed connections on servers
- Data sharing options do exist for customers willing to accept risk
- Shared disk resources can be setup between VMs
  - VMM shared virtual disks
  - Shared disk partitions
  - Shared disk can be ACLed to prevent writes from higher levels

Using Single NIC & VPNs

- Exploring use of single cable and VPNs to access these different domains
  - Used commonly in thin client architectures
- Using VPNs to access low networks over system high DMZ seems promising
  - IPSec policy on each guest OS
    - “from me to any use gateway *** require AH with cert foo and ESP”
  - Each guest OS would maintain a cert and keys that can only access one network VPN server
    - Certs issued by CA on respective network
    - Same layers of isolation protect these keys from other VMs
  - All traffic on DMZ is strong encrypted
    - 3DES keyed with 2048-bit DH supported in Windows IPS

Using Single NIC & VPNs

- VPN servers allow connection from authenticated guest VM to back end network resources
- VPN servers ensures correct VM authenticated and proper encryption used
  - Assume the guest VM is compromised and the server will still prevent it from connecting to the wrong network
- Hybrid approach is to use one NIC per level
  - NIC 1 connects to multiple TS networks
  - NIC 2 connects to multiple Secret networks

Typhon VPN Solution

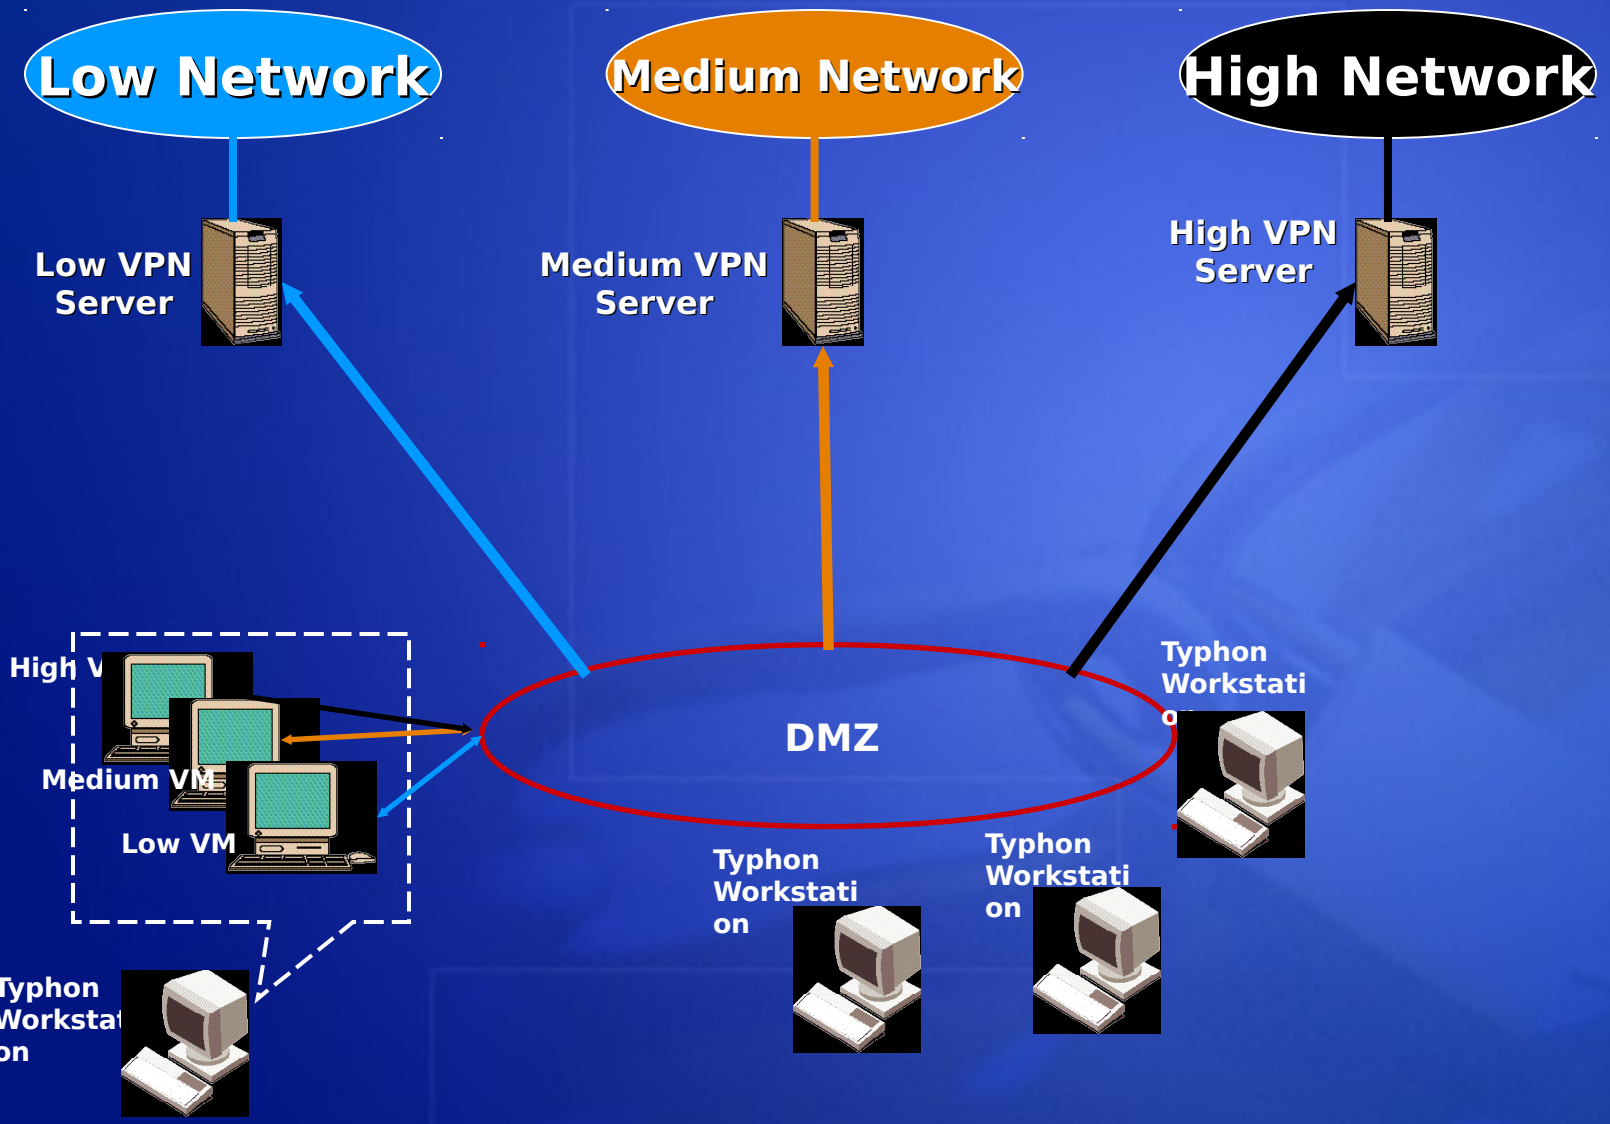

More Demonstration

The Typhon XPerience

Trusted Multi-Net: HatWizard

A Configurable Thin Client Architecture to Access Multiple Security Domains

Bill Neumann
MCS Principal Security Architect
Microsoft Federal
wneumann@microsoft.com

MITRE TECHNOLOGY APPLIED

Microsoft Windows

Background

- Trusted Multi-Net: HatWizard based on the Unified Cryptologic Architecture Office's "HatWizard" project—architected by MITRE
  - Four years of development, refinement, and testing
  - Two rounds of security testing
  - One operational pilot
  - Microsoft partnering with Arrowhead Global Solutions, Citrix, and Concurrent Technologies Corporation
- HatWizard based on Win XPe thin clients, server-based computing, a smart-card based PKI, and VPN tunneling

HatWizard Problem

User And IT Needs

- Secure desktop access to information
  - On any network
  - At any security classification level
  - From anywhere
- Secure real-time and asynchronous collaboration
  - To anyone
  - From anywhere
- Establish ad hoc classified networks to support coalition operations
- Integrate secure access and collaboration with existing IT network infrastructure

HatWizard's Big Issues

- How do you deliver the bits to the desktop from multiple networks operating at different security classification levels?
- How do you secure a desktop device that has access to information at different security classification levels?
- How do you provide users with secure real-time and asynchronous collaboration over networks operating at different security classification levels?
- How do you establish an ad hoc classified network environment?

Trusted Multi-Net: HatWizard
Provides Security and Coexistence

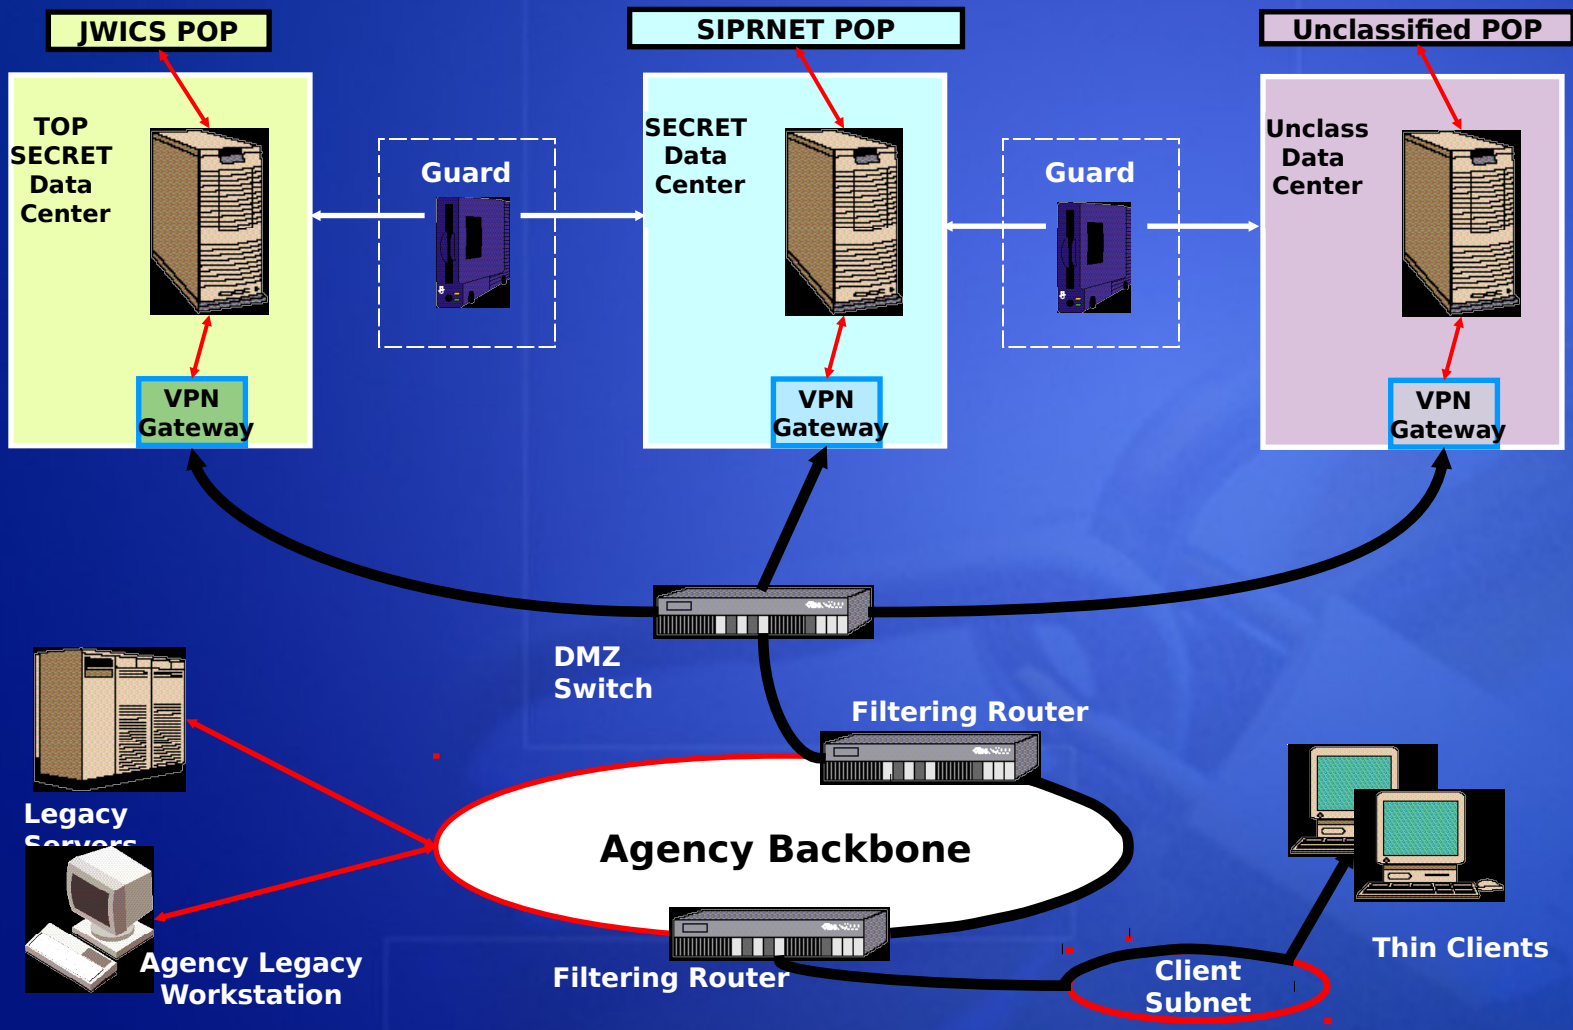

Win XPe Thin Client Desktop

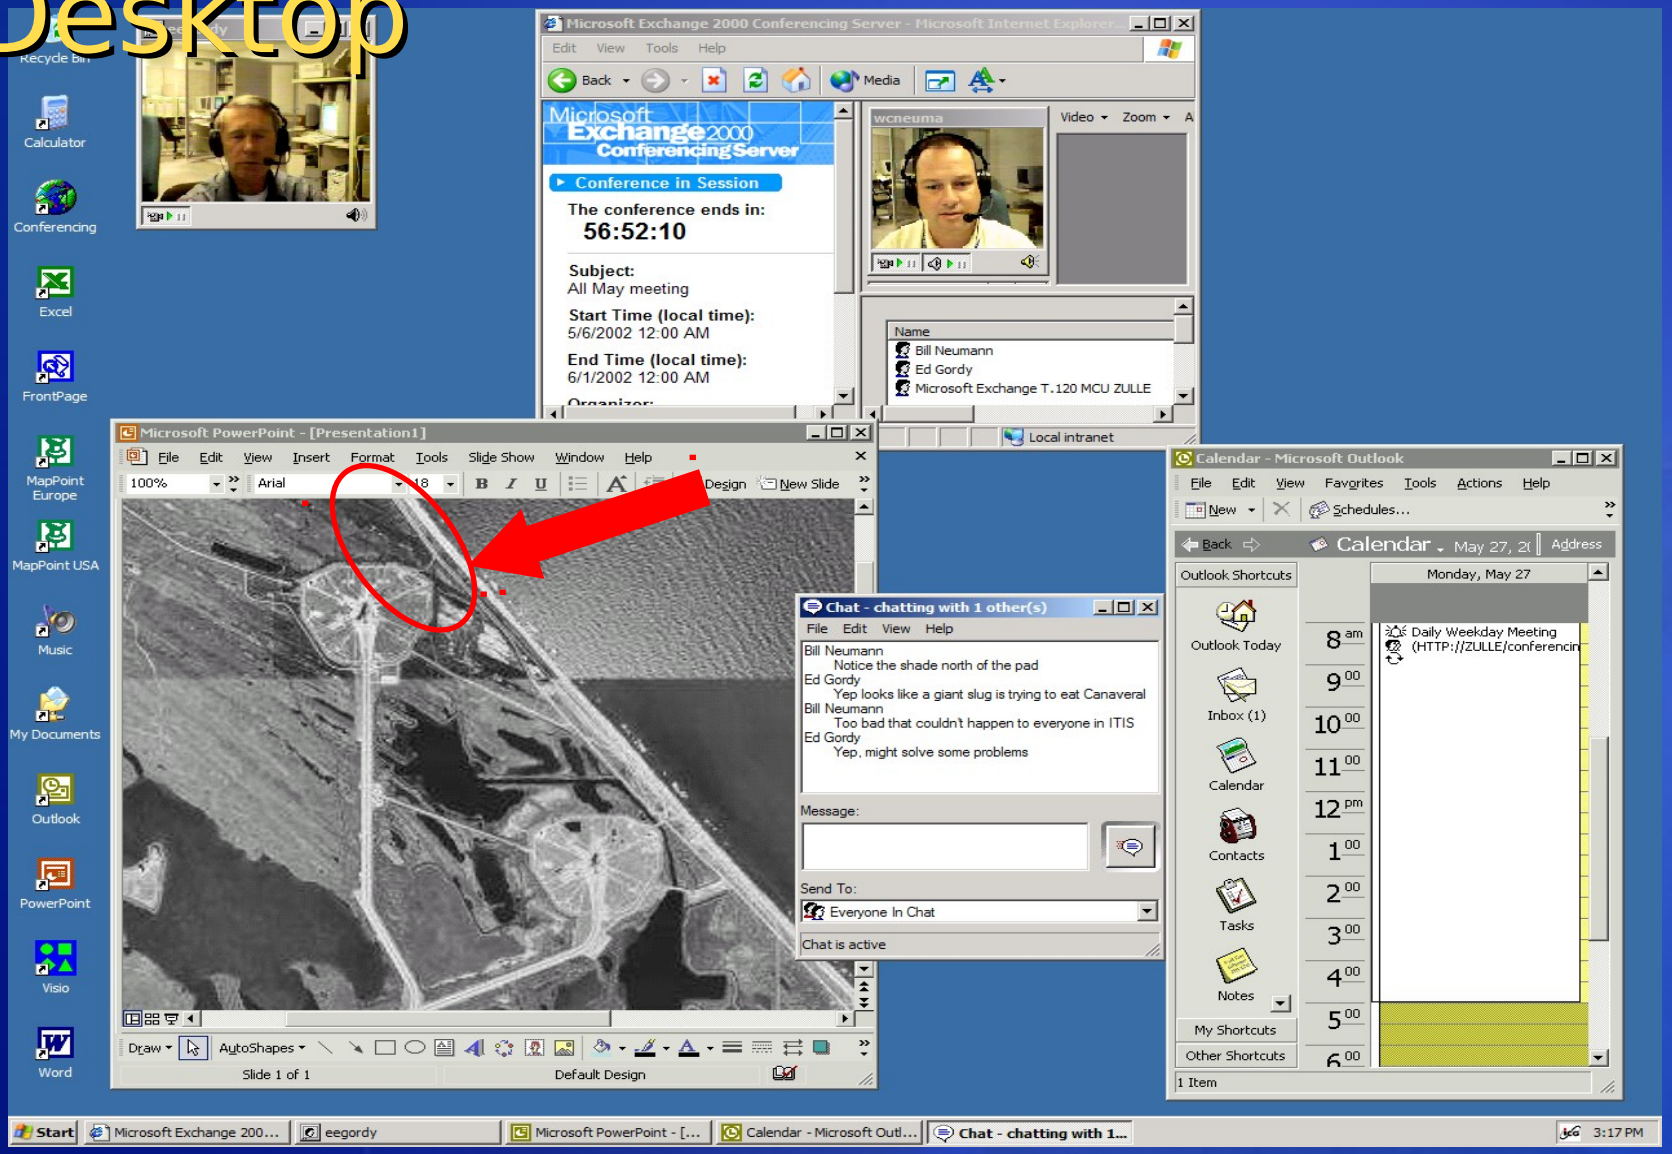

The NYTOR Raven

- No hard disk, floppy disk, or CD-ROM
- Win XPe run in locked down Flash RAM
- Win XPe image configured for secure operations
- No local user logon available!

Trusted Multi-Net: HatWizard
Primary Benefits

- Provides *secure serial connectivity to an unlimited number of network environments*
- *Real-time audio and video collaboration is supported* through a resident NetMeeting application that runs within a locked down Internet Explorer browser on the thin client
- *Both users and the thin clients are strongly authenticated* using digital certificates as part of establishing VPN tunnels and logging onto any network environment
- Can also be applied to health care, banking, insurance, and software development verticals
  - Secure remote access to applications and data
  - Rigorous control of applications, information, and storage on the desktop device
  - Strong machine and user authentication

Trusted Multi-Net: HatWizard
Additional Benefits

- Flexible support for coalition operations
- Increased bandwidth efficiency
- Applications integration
- Wireless network capable
- Greater resistance to insider attacks/ malicious code
- Integrated network and systems management
- Easier configuration management of software applications

Trusted Multi-Net: HatWizard Limitations

- Two rounds of NSA security testing identified some issues
  - Incremental technical improvements have been incorporated into the solution—no architectural issues have been identified
- Initial security domain configuration is complex
  - New Group Policy Console simplifies management of group Management policies
- Migration of legacy desktop application environments can be costly and time consuming
  - Blade computers can replace or augment server-based computers
- Currently supports serial VPN connections
  - New NYTOR Eclipse thin client uses a Typhon XP-like VMM and 3 smart card readers to support simultaneous VPN connections

Trusted Multi-Net: HatWizard
Status

- DIA evaluating 3 different MSL desktop solutions one of which is Trusted Multi-Net
  - Single VPN and VMM variants
  - Use of blades and server-based computing for applications processing
- New NYTOR Eclipse thin client solution can be seen in Partner Pavilion

Questions?